

A Java Implementation of Parameter-less Evolutionary Algorithms

José C. Pereira

CENSE and DEEI-FCT
Universidade do Algarve
Campus de Gambelas
8005-139 Faro, Portugal
unidadeimaginaria@gmail.com

Fernando G. Lobo

CENSE and DEEI-FCT
Universidade do Algarve
Campus de Gambelas
8005-139 Faro, Portugal
fernando.lobo@gmail.com

Abstract

The Parameter-less Genetic Algorithm was first presented by Harik and Lobo in 1999 as an alternative to the usual trial-and-error method of finding, for each given problem, an acceptable set-up of the parameter values of the genetic algorithm. Since then, the same strategy has been successfully applied to create parameter-less versions of other population-based search algorithms such as the Extended Compact Genetic Algorithm and the Hierarchical Bayesian Optimization Algorithm. This report describes a Java implementation, *Parameter-less Evolutionary Algorithm* (P-EAJava), that integrates several parameter-less evolutionary algorithms into a single platform. Along with a brief description of P-EAJava, we also provide detailed instructions on how to use it, how to implement new problems, and how to generate new parameter-less versions of evolutionary algorithms.

At present time, P-EAJava already includes parameter-less versions of the Simple Genetic Algorithm, the Extended Compact Genetic Algorithm, the Univariate Marginal Distribution Algorithm, and the Hierarchical Bayesian Optimization Algorithm. The source and binary files of the Java implementation of P-EAJava are available for free download at <https://github.com/JoseCPereira/2015ParameterlessEvolutionaryAlgorithmsJava>.

1 Introduction

The parameter-less genetic algorithm (P-GA) was first presented by Harik and Lobo (1999) as an alternative to the usual trial-and-error method of finding, for each given problem, an acceptable set-up of the parameter values of the genetic algorithm.

Shortly after, Pelikan and Lobo (1999) showed that, even in the worst-case scenario, the total number of function evaluations performed in one run of the P-GA does not increase significantly with respect to the number of function evaluations performed by the standalone GA tuned with an optimal fixed population size. Moreover, the authors pointed out that the worst-case scenario was extremely improbable to occur and that the expected performance of the P-GA in practice should

be much more benevolent. Therefore, Pelikan and Lobo (1999) concluded that the P-GA is an efficient method of eliminating the parameters of the simple GA .

Since then, the same strategy has been successfully applied to develop other parameter-less evolutionary algorithms such as the Parameter-less Extended Compact Genetic Algorithm (P-ECGA) (Lobo, 2000), the Parameter-less Genetic Programming Algorithm (P-GPA) (Spinosa and Pozo, 2002), and the Parameter-less Hierarchical Bayesian Optimization Algorithm (P-HBOA) (Pelikan et al., 2007). Building on these results, we implemented in Java the *Parameter-less Evolutionary Algorithm* (P-EAJava) that integrates several parameter-less evolutionary algorithms into a single platform.

At present time, P-EAJava already includes the P-GA, the P-ECGA, the P-HBOA, and a parameter-less version of the Univariate Marginal Distribution Algorithm (Mühlenbein and Paaß, 1996). The source and binary files of the Java implementation of P-EAJava are available for free download at <https://github.com/JoseCPereira/2015ParameterlessEvolutionaryAlgorithmsJava>.

The P-EAJava uses as a common working base the standard versions of the Simple Genetic Algorithm (Holland, 1975; Goldberg, 1989; Eiben and Smith, 2003), the Univariate Marginal Distribution Algorithm (Mühlenbein and Paaß, 1996), the Extended Compact Genetic Algorithm (Harik, 1999), and the Hierarchical Bayesian Optimization Algorithm (Pelikan and Goldberg, 2006). The Java implementation of these four standard EAs is presented in detail in another arXiv report from the same authors. The corresponding source code is also available for free download at <https://github.com/JoseCPereira/2015EvolutionaryAlgorithmsJava>.

The remainder of this paper is organized as follows. In Section 2 we briefly describe the main concepts of the P-GA that are the basis for the P-EAJava implementation. In Section 3 we discuss the P-EAJava implementation itself and provide detailed instructions on how to use it, how to implement new problems with it, and how to generate new parameter-less versions of evolutionary algorithms.

2 The Parameter-less Genetic Algorithm

The P-GA establishes an evolutionary race among an unbounded number of independent populations. All populations are evolved using the same simple GA which has most of its parameters set a priori to robust fixed values, in accordance with facetwise theoretical studies (Goldberg et al., 1993; Thierens and Goldberg, 1993). The only exception is the population size which is dynamically adapted using a population sizing method.

In the P-GA each independent population is uniquely identified by its size, N_i . This size doubles between consecutive populations, according to the identity

$$N_i = 2^i N_0, \quad \forall i \in \mathbb{N} . \quad (1)$$

where N_0 is the initial population size, chosen to be a small enough value such as $N_0 = 10$.

The P-GA starts by performing 4 generations of population N_0 , after which it evolves population N_1 for a single generation. This process is then repeated so that, for each 4 generations performed with population N_0 , there will be one generation performed with population N_1 . The same evolutionary procedure is used recursively for all other populations so that, throughout the entire run of the algorithm, for each 4 generations performed with population N_i , there will be

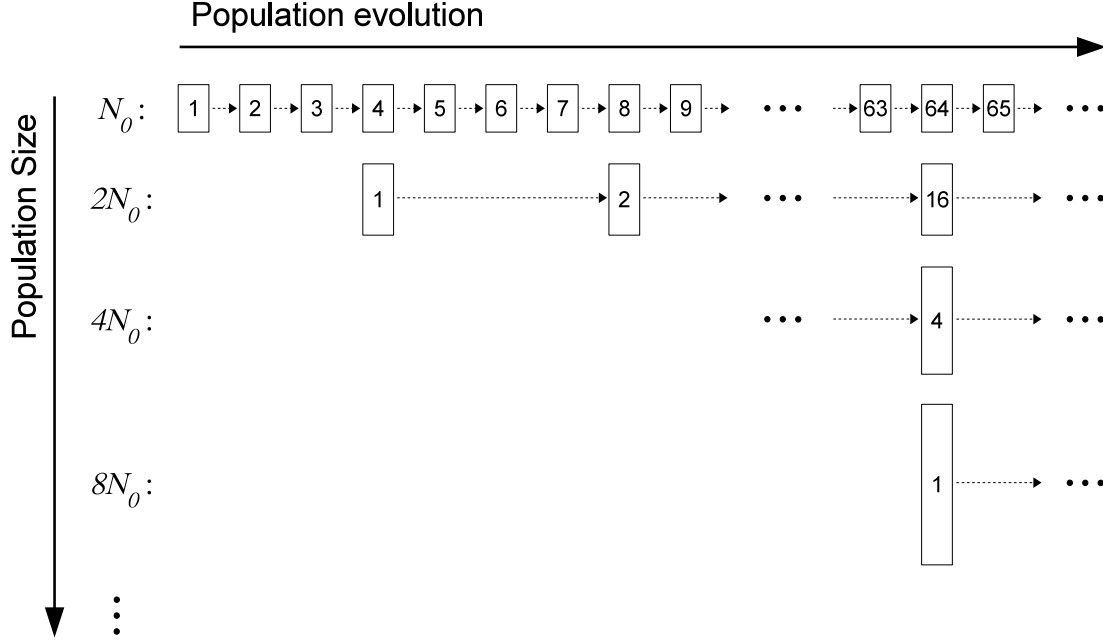


Figure 1: The P-GA establishes a race among an unbounded number of populations with exponentially increasing size that are evolved using a simple GA. For each 4 generations performed with population $2^i N_0$ there will be one generation performed with population $2^{i+1} N_0$. The initial population size N_0 is chosen to be a small enough value such as $N_0 = 10$.

one generation performed with population N_{i+1} . Figure 1 depicts a graphical representation of a possible instance of this population evolution.

The P-GA works under the assumption that solution quality grows monotonously with the population size (Lobo and Lima, 2007). In addition, the algorithm also assumes that this improvement in solution quality happens at a decreasing rate, i.e., when the same increase in population size happens between larger and larger populations, the gains in solution quality are expected to be fewer and fewer. For these reasons, the evolutionary race is established between populations with *exponentially increasing* size. Simultaneously, the P-GA allows smaller populations to use more computational resources, because if two different populations produce solutions with similar quality then the smaller population should be preferred, all other things being equal.

In their paper, Harik and Lobo (1999) used a counter of base 4 to implement the population sizing method of the P-GA (see also, Lobo (2000); Pelikan and Lobo (1999) for detailed descriptions of this counter). However, Pelikan et al. (2007) suggested a somewhat simpler implementation for that sizing method and used it to propose a parameter-less version of the HBOA. The same simpler implementation was adopted to develop the work reported in this paper. Figure 2 depicts the pseudocode for the population sizing method as presented by Pelikan et al, where only the name in the title had to be changed to make it work for the simple genetic algorithm.

In the P-GA, the performance of a population at any point in time is measured by the current average fitness of all its individuals and by how many resources it has consumed, i.e., by how many fitness function evaluations it has already performed. Naturally, a population that has consumed more resources is expected to have a better fitness than any other population that has performed

P-GA POPULATION SIZING METHOD

```

1  initialize population[0] with size  $N_0$ 
2  generation[0]  $\leftarrow 0$ 
3  max_initialized  $\leftarrow 0$ 
4  i  $\leftarrow 0$ 
5  while (not done)
6      do evolve population[i] one generation
7          generation[i]++
8          if (generation[i] mod 4 = 0)
9              then i++
10             if (i > max_initialized)
11                 then initialize population[i] with size  $2^i N_0$ 
12                     max_initialized  $\leftarrow i$ 
13             else i  $\leftarrow 0$ 

```

Figure 2: Pseudocode for the population sizing method employed by the P-GA.

less fitness evaluations. When this is not the case, the more costly populations should be discarded. In practice, whenever population N_i obtains an average fitness at least as good as the average fitness of population N_{i-1} , the P-GA eliminates all smaller populations N_k , with $k < i$. Note that, strictly speaking, the P-GA could maintain active those smaller populations that eventually still have a better fitness than N_i . However, if N_i has already caught up with N_{i-1} , it will definitely catch up with all other smaller populations, rather sooner than later. Therefore, it is more efficient not to waste any more time with these populations and focus all resources on those populations that promise to deliver better quality solutions.

For the sake of simplicity the P-GA employs a genetic algorithm without a mutation operator. In the absence of mutation, all individuals in a population eventually become identical. In such a state, the population is said to have converged and can no longer generate any new individuals and will not improve its average fitness. In this case, the P-GA stores the information concerning the best solution found and eliminates not only the converged population, but all other smaller populations. This is efficient and effective, because once a population has converged, all smaller populations will also converge in the near future, and since solution quality increases along side with population size, there is nothing to gain by waiting for them to do so.

The P-GA is designed to run forever, physical constraints aside, because in practice the quality of the optimal solution is often unknown for many problems, making it impossible to distinguish, for instance, when the algorithm has reached the optimum result from when it simply got “stuck” in some plateau of the search space. Therefore, the P-GA leaves to the user the decision when to stop the computation, based on the quality of the solutions already found and on the time and resources that she or he is willing to spend.

With an ever increasing set of larger and larger populations, the P-GA would soon exhaust all computational resources available. By discarding populations that either do not comply with the parameter-less invariant or can no longer improve the quality of its solutions, the P-GA ensures that the actual number of active populations, although unbounded, will not grow excessively large

with time. In fact, this discarding process, which is more intensive for smaller populations, is one of the main factors behind the relatively small amount of extra computational resources that the P-GA requires in relation to the standalone GA working with an optimal fixed population size.

The implementation of a parameter-less version of the HBOA by Pelikan et al. (2007) showed in practice that the parameter-less strategy proposed by Harik and Lobo (1999) was not tied to a specific evolutionary algorithm. The same result had also been shown by Lobo (2000) with the implementation of the P-ECGA which was used by the author to solve an electrical network expansion problem. Likewise, Pelikan et al. (2007) used the P-HBOA to successfully solve artificial hierarchical and nearly decomposable problems, and a 2D version of the real world Ising Spin Glasses problem. In the next section, we describe how to use the P-EAJava.

3 The Parameter-less Evolutionary Algorithm

The Parameter-less Evolutionary Algorithm (P-EAJava) is a Java implementation of the population sizing method employed by the P-GA (see Figure 2). However, by decoupling the population sizing method from the evolutionary algorithm itself, the P-EAJava allows the creation of parameter-less versions of other evolutionary algorithms in a straightforward manner.

Naturally, the P-EAJava works with algorithms for which it is reasonable to presume the same assumptions as the simple GA, namely that solution quality grows monotonously with the population size, but at a decreasing rate and that it is possible to effectively automate the adaptation process of all the algorithm's parameters, with the possible exception of the population size.

At present time, an evolutionary algorithm must satisfy the following set of constraints in order to be integrated in the P-EAJava:

1. The algorithm represents possible solutions (individuals) as strings of zeros and ones.
2. All individuals have the same string size.
3. The population size remains constant throughout a complete run of the algorithm.

Naturally, it should be possible to further generalize the P-EAJava and eliminate or at least weaken such constraints.

3.1 How to use the P-EAJava

The P-EA is a Java application developed with the Eclipse¹ IDE. The available code is already compiled and can be executed using the command line.

Run the P-EAJava from a command line

1. Unzip the source file *2015ParameterlessEAs.zip* to any directory.

¹Version: Kepler Service Release 2

2. Open your favourite terminal and execute the command

```
cd [yourDirectory]/2015ParameterlessEvolutionaryAlgorithmsJava/bin
```

where [yourDirectory] is the name of the directory chosen in step 1.

3. Execute the command

```
java com/z_PEA/PARAMETERLESS ./PEAParameters.txt
```

The argument “PEAParameters.txt” is in fact the name of the file containing all the options concerning the P-EA settings and can be changed at will.

After each execution of a single or multiple runs, the P-EAJava produces one output file – *PARAMETERLESS_*.txt* – that records how each run progressed in terms of population size, best current fitness, among other relevant information. Additionally, the P-EAJava also creates the file – *PARAMETERLESS-STATS_*.txt* – that stores some of the statistics necessary for analyzing the behaviour of the parameter-less algorithms over multiple runs.

At present time, the P-EAJava version made available with this paper already includes four parameter-less algorithms:

- Parameter-less Genetic Algorithm,
- Parameter-less Univariate Marginal Distribution Algorithm,
- Parameter-less Extended Compact Genetic Algorithm,
- Parameter-less Hierarchical Bayesian Optimization Algorithm

The current code also includes a set of test problems that can be solved using some or all of the previous algorithms. Here is the problem menu:

ZERO Problems

- 0 → ZeroMax
- 1 → Zero Quadratic
- 2 → Zero 3-Deceptive
- 3 → Zero 3-Deceptive Bipolar
- 4 → Zero 3-Deceptive Overlapping
- 5 → Zero Concatenated Trap-k
- 6 → Zero Uniform 6-Blocks

ONE Problems

- 10 → OneMax
- 11 → Quadratic
- 12 → 3-Deceptive
- 13 → 3-Deceptive Bipolar
- 14 → 3-Deceptive Overlapping
- 15 → Concatenated Trap-k
- 16 → Uniform 6-Blocks

Hierarchical Problems

- 21 → Hierarchical Trap One
- 22 → Hierarchical Trap Two

The Zero problems always have the string with all zeros as their best individual. The One problems are the same as the Zero problems but their best individual is now the string with all ones. A description of these problems can be found, for instance, in Pelikan et al. (2000). The Hierarchical problems are thoroughly described in Pelikan (2005).

It is also possible to define a noisy version for any of the previous problems. This is done by adding a non-zero Gaussian noise term to the fitness function.

The source code that implements all the problems mentioned in this section can be found in the file *src/com/z_PEA/Problem.java*.

As mentioned previously, all options concerning the parameter-less strategy are in the file *PEAParameters.txt*. In particular, it is in this file that are made the choices for the problem to be solved and for the parameter-less algorithm that is going to solve the chosen problem.

To choose a particular problem the user must set the value of the following three options:

Line 81: *problemType*
Line 90: *stringSize*
Line 107: *sigmaK* (defines the noise component)

Analogously, to choose a particular parameter-less algorithm it is necessary to set the values of the following two options:

Line 125: *eAlg*
Line 135: *eaParamFile* (defines the name of the corresponding parameters file)

All other options are set to default values and their role in the parameter-less strategy is explained with detail in the file's comments. This is also true for the parameters specific to each of the implemented algorithms which are defined in four separate files:

SGA: *SGAParameters.txt*
UMDA: *UMDAParameters.txt*
ECGA: *ECGAParameters.txt*
HBOA: *HBOAParameters.txt*

Note that the default settings defined in these three files were chosen to ensure a robust behavior of the corresponding algorithms, in accordance with current theory. Therefore, the user is advised to proceed with caution when performing any changes in those settings. In fact, the whole idea behind the parameter-less strategy is to eliminate the need of such fine tuning when solving a particular problem. After choosing a problem to be solved and a particular algorithm to solve it, the user has only to press the start button and wait until the P-EAJava finds a solution with good enough quality.

3.2 How to implement a new problem with P-EAJava

The P-EAJava uses the design pattern strategy (Gamma et al., 1995) to decouple the implementation of a particular problem from the remaining parameter-less strategy (see Figure 3). As a consequence, to plug in a new problem to the framework it is only necessary to define one class that implements the interface *IProblem* and change some input options to include the new choice. The interface *IProblem* can be found in the file *src/com/z_PEA/Problem.java*.

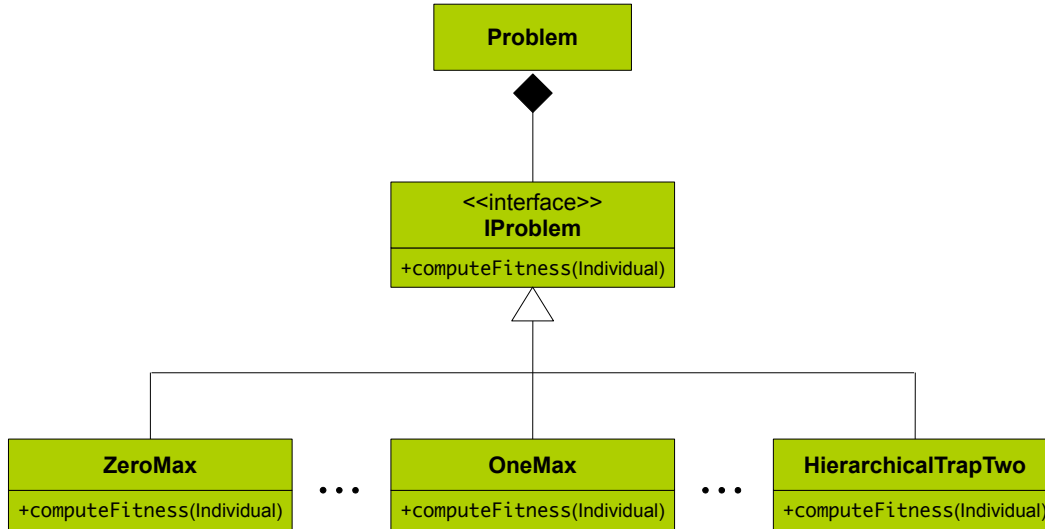


Figure 3: The P-EAJava uses the design pattern strategy (Gamma et al., 1995) to allow an easy implementation of new problems to be solved by the framework.

In the following let us consider that we want to solve a new problem called *NewProblem* with one of the parameter-less algorithms. To plug in this problem it is necessary to:

1. Define a class called *NewProblem* in the file *src/com/z_PEA/Problem.java*. The signature of the class will be

```

class NewProblem implements IProblem{
    public float computeFitness(Individual individual){
        // TODO
    }
    public String toString(){
        return "NEW PROBLEM";
    }
}

```

2. Code the body of the function `computeFitness(Individual)` according to the nature of problem *newProblem*. The class *Individual* provides all the necessary functionalities to operate with the string of zeros and ones that represents an individual (e.g., `getAllele(int)`). This class can be found in the file *src/com/z_PEA/Individual.java*.
3. To define the new problem option, add the line

```

case 99: return new Problem(new NewProblem(), stringSize, optimumValue, sigma);

```

to the **switch** command in line 209 of the file *src/com/z_PEA/ParParameter.java*. The case number – 99 – is a mere identifier of the new problem option. The user is free to choose other value for this purpose. The rest of the line is to be written verbatim.

4. Validate the new problem option value – 99 – by adding the case `problemType == 99` to the conditional in line 105 of the same *ParParameter.java* file.

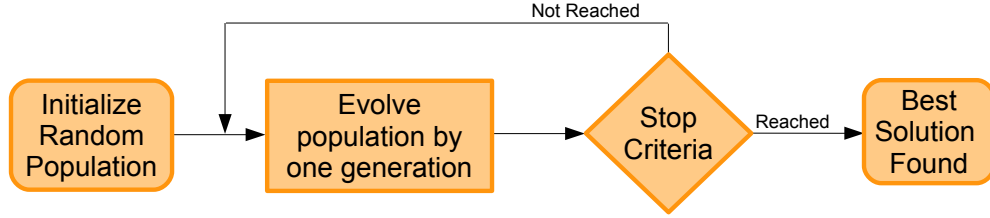


Figure 4: General flowchart of a population-based search algorithm. The ParFrame can only integrate algorithms that follow this workflow.

Although not strictly necessary, it is also advisable to keep updated the problem menu in the file *PEAParameters.txt*.

3.3 How to create a new parameter-less algorithm with P-EAJava

At present time, the P-EAJava is general enough to work with population-based algorithms that follow the workflow depicted in Figure 4. As such, the framework already provides all the necessary classes that define the working notions of Population and Stop Criteria. The evolutionary algorithm has only to provide its own classes related with the operators responsible for the evolution of the population in each generation.

The P-EAJava also uses the design pattern strategy (Gamma et al., 1995) to decouple as much as possible the implementation of a search algorithm from the remaining parameter-less strategy (see Figure 5). However, as expected, to create a new parameter-less algorithm is not as simple as plugging in a new problem.

Let us look at the example of the parameter-less SGA to better understand all the necessary steps. Note that, the reading of the following instructions is best complemented with a good analysis of the corresponding source code.

1. Define the class *SGA* that implements the interface *IEAlgorithm*. Here is the important part of this implementation:

```

class SGA extends IEAlgorithm{
    //...
    public IEASolver newIEASolver(int currentN){
        return new SGASolver(paramFile, currentN);
    }
    //..
}

```

Note that, the function *newIEASolver(int)* has the return type *IEASolver* which is itself another interface. The class *SGA* is included in the file *src/com/z_PEA/SGA.java*.

2. Define the class *SGASolver* that implements the interface *IEASolver*. This class includes the all-important function *nextGeneration()*, responsible for implementing the evolution of the population in each generation. Here is the main part of the implementation:

```

public boolean nextGeneration(){

    SelectedSet selectedSet = selection.select(currentPopulation);
    Individual[] newIndividuals = crossover.cross(selectedSet);
    mutation.mutate(newIndividuals);
    replacement.replace(currentPopulation, newIndividuals);

    currentGeneration++;
    fitnessCalls += newIndividuals.length;
    currentPopulation.computeUnivariateFrequencies();
    avgFitness = currentPopulation.computeAvgFitness();

    return Stopper.criteria(currentGeneration, currentPopulation);
}

```

Note that, the function must return a boolean as the result of evaluating the stop criteria. The class *SGASolver* is included in the file *src/com/SGA/SGASolver.java*.

3. To define the new algorithm option, add the line

```
case 0: return new SGA(eaParamFile);
```

to the **switch** command in line 241 of the file *src/com/z_PEA/ParParameter.java*.

4. Validate the new problem option value – 0 – by adding the case *eAlg == 0* to the conditional in line 139 of the same *ParParameter.java* file.

In order to achieve complete integration in the P-EAJava, the parameter-less SGA must also use the classes provided by the framework that define, for instance, the stop criteria of the algorithm. Here is a brief description of those classes:

Population Contains the array of individuals that are the population. It provides all the functionalities necessary for operating that set of individuals. In particular, it is responsible for computing its own average fitness, which is a crucial information for the parameter-less strategy.

RandomPopulation Subclass of the class *Population*. The constructor of this class is responsible for generating the initial random population.

SelectedSet Subclass of the class *Population*. For some algorithms, depending on the selection operator, the size of this set is different from the population size.

Individual Contains the string of zeros and ones that represents an individual. It provides all the functionalities necessary for operating that string. In particular, it is responsible for computing its own fitness, according to the problem at hand.

Stopper Contains all the stop criteria used by the parameter-less algorithms. These criteria are integrated in a single function called *criteria(...)* which in turn must be returned by the *nextGeneration()* function (see step 2 of the SGA instructions). The criteria options can be changed in the file *PEAParameters.txt*

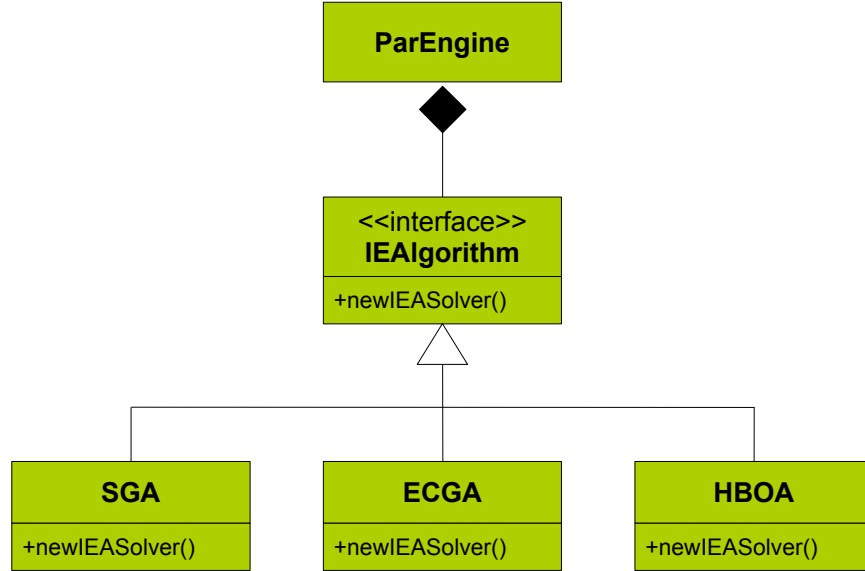


Figure 5: The P-EAJava uses the design pattern strategy (Gamma et al., 1995) to decouple as much as possible the implementation of a search algorithm from the remaining parameter-less strategy.

Acknowledgements

The current version of the P-EAJava is one of the by-products of the activities performed by the PhD fellow, José C. Pereira within the doctoral program entitled *Search and Optimization with Adaptive Selection of Evolutionary Algorithms*. The work program is supported by the Portuguese Foundation for Science and Technology with the doctoral scholarship SFRH/BD/78382/2011 and with the research project PTDC/EEI-AUT/2844/2012.

References

- Eiben, A. E. and Smith, J. E. (2003). *Introduction to Evolutionary Computing*. SpringerVerlag.
- Gamma, E., Helm, R., Johnson, R., and Vlissides, J. (1995). *Design patterns: elements of reusable object-oriented software*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA.
- Goldberg, D. E. (1989). *Genetic algorithms in search, optimization, and machine learning*. Addison-Wesley, Reading, MA.
- Goldberg, D. E., Deb, K., and Thierens, D. (1993). Toward a better understanding of mixing in genetic algorithms. *Journal of the Society of Instrument and Control Engineers*, 32(1):10–16.
- Harik, G. R. (1999). Linkage learning via probabilistic modeling in the ECGA. IlliGAL Report No. 99010, Illinois Genetic Algorithms Laboratory, University of Illinois at Urbana-Champaign, Urbana, IL.
- Harik, G. R. and Lobo, F. G. (1999). A parameter-less genetic algorithm. In Banzhaf, W. et al.,

- editors, *Proceedings of the Genetic and Evolutionary Computation Conference GECCO-99*, pages 258–265, San Francisco, CA. Morgan Kaufmann.
- Holland, J. H. (1975). *Adaptation in Natural and Artificial Systems*. University of Michigan Press, Ann Arbor, MI.
- Lobo, F. G. (2000). *The parameter-less genetic algorithm: Rational and automated parameter selection for simplified genetic algorithm operation*. PhD thesis, Universidade Nova de Lisboa, Portugal. Also IlliGAL Report No. 2000030.
- Lobo, F. G. and Lima, C. F. (2007). Adaptive population sizing schemes in genetic algorithms. In *Parameter Setting in Evolutionary Algorithms*, pages 185–204. Springer.
- Mühlenbein, H. and Paaß, G. (1996). From recombination of genes to the estimation of distributions I. Binary parameters. In Voigt, H.-M. et al., editors, *Parallel Problem Solving from Nature – PPSN IV*, pages 178–187, Berlin. Kluwer Academic Publishers.
- Pelikan, M. (2005). *Hierarchical Bayesian Optimization Algorithm: Toward a New Generation of Evolutionary Algorithms*. Springer.
- Pelikan, M. and Goldberg, D. (2006). Hierarchical bayesian optimization algorithm. In Pelikan, M., Sastry, K., and CantúPaz, E., editors, *Scalable Optimization via Probabilistic Modeling*, volume 33 of *Studies in Computational Intelligence*, pages 63–90. Springer Berlin Heidelberg.
- Pelikan, M., Goldberg, D. E., and Cantú-Paz, E. (2000). Linkage problem, distribution estimation, and Bayesian networks. *Evolutionary Computation*, 8(3):311–341.
- Pelikan, M., Hartmann, A., and Lin, T.-K. (2007). Parameter-less hierarchical bayesian optimization algorithm. In *Parameter Setting in Evolutionary Algorithms*, volume 54, pages 225–239. Springer Berlin Heidelberg, Berlin.
- Pelikan, M. and Lobo, F. G. (1999). Parameter-less genetic algorithm: A worst-case time and space complexity analysis. IlliGAL Report No. 99014, University of Illinois at Urbana-Champaign, Illinois Genetic Algorithms Laboratory, Urbana, IL.
- Spinosa, E. and Pozo, A. (2002). Controlling the population size in genetic programming. In *Proceedings of the 16th Brazilian Symposium on Artificial Intelligence: Advances in Artificial Intelligence*, SBIA '02, pages 345–354, London, UK, UK. Springer-Verlag.
- Thierens, D. and Goldberg, D. E. (1993). Mixing in genetic algorithms. In Forrest, S., editor, *Proceedings of the Fifth International Conference on Genetic Algorithms*, pages 38–45, San Mateo, CA. Morgan Kaufmann.